A Note on Easy and Efficient Computation of Full Abelian Periods of a Word[1]

G. Fici[a], T. Lecroq[b], A. Lefebvre[b], É. Prieur-Gaston[b], W. F. Smyth[c]

[a] *Dipartimento di Matematica e Informatica, Università di Palermo, Italy*
[b] *LITIS EA4108, Université de Rouen, 76821 Mont-Saint-Aignan Cedex, France*
[c] *Dept of Computing and Software, McMaster University, Hamilton ON L8S 4K1, Canada and Faculty of Engineering & Information Technology, Murdoch University, Murdoch WA 6150, Australia*

arXiv:1510.00634v1 [cs.DS] 2 Oct 2015

Abstract

Constantinescu and Ilie (Bulletin of the EATCS 89, 167–170, 2006) introduced the idea of an Abelian period with head and tail of a finite word. An Abelian period is called full if both the head and the tail are empty. We present a simple and easy-to-implement $O(n \log \log n)$-time algorithm for computing all the full Abelian periods of a word of length n over a constant-size alphabet. Experiments show that our algorithm significantly outperforms the $O(n)$ algorithm proposed by Kociumaka et al. (Proc. of STACS, 245–256, 2013) for the same problem.

Keywords: Abelian period; Abelian power; weak repetition; design of algorithms; text algorithms; combinatorics on words.

1. Introduction

The study of repetitions in words is a classical topic in Stringology. A word is called an (integer) power if it can be written as the concatenation of two or more copies of another word, like `barbar`. However, any word can be written as a *fractional* power; that is, given a word $\boldsymbol{w}$, one can always find a word $\boldsymbol{u}$ such that $\boldsymbol{w} = \boldsymbol{u}^n\boldsymbol{u'}$, where $\boldsymbol{u'}$ is a (possible empty) prefix of $\boldsymbol{u}$ and n is an integer greater than or equal to one. In this case, the length of $\boldsymbol{u}$ is called *a period* of the word $\boldsymbol{w}$. A word $\boldsymbol{w}$ can have different periods, the least of which is usually called *the period* of $\boldsymbol{w}$.

Recently, a natural extension of this setting has been considered involving the notion of commutative equivalence. Two words are called commutatively equivalent if they have the same number of occurrences of each letter; that is, if one is an anagram of the other. An Abelian power (also called a weak repetition [2]) is a word that can be written as the concatenation of two or more words that are commutatively equivalent, like `iceddice`.

Recall that the Parikh vector $\mathcal{P}_{\boldsymbol{w}}$ of a word $\boldsymbol{w}$ is the vector whose ith entry is the number of occurrences of the ith letter of the alphabet in $\boldsymbol{w}$. For example, given the (ordered) alphabet $\Sigma = \{a, b, c\}$, the Parikh vector of the word $\boldsymbol{w} = aaba$ is $\mathcal{P}_{\boldsymbol{w}} = (3, 1, 0)$. Two words are therefore commutatively equivalent if and only if they have the same Parikh vector.

Constantinescu and Ilie [3] introduced the definition of an Abelian period with head and tail of a word $\boldsymbol{w}$ over a finite ordered alphabet $\Sigma = \{a_1, a_2, \ldots, a_\sigma\}$: An integer $p > 0$ is an Abelian period of $\boldsymbol{w}$ if one can write $\boldsymbol{w} = \boldsymbol{u}_0\boldsymbol{u}_1 \cdots \boldsymbol{u}_{k-1}\boldsymbol{u}_k$ where for $0 < i < k$ all the factors $\boldsymbol{u}_i$'s have the same Parikh vector $\mathcal{P}$ such that $\sum_{j=1}^{\sigma} \mathcal{P}[j] = p$ and the Parikh vectors of $\boldsymbol{u}_0$ and $\boldsymbol{u}_k$ are "contained" in $\mathcal{P}$, in the sense that they are

[1] The results in this note have been presented in preliminary form in [1].

Email addresses: `Gabriele.Fici@unipa.it` (G. Fici), `Thierry.Lecroq@univ-rouen.fr` (T. Lecroq), `Arnaud.Lefebvre@univ-rouen.fr` (A. Lefebvre), `Elise.Prieur@univ-rouen.fr` (É. Prieur-Gaston), `smyth@mcmaster.ca` (W. F. Smyth)

proper sub-Parikh vectors of $\mathcal{P}$ (see next section for the formal definition of "contained"). In this case, $\boldsymbol{u}_0$ and $\boldsymbol{u}_k$ are called the head and the tail of the Abelian period p, respectively. This definition of an Abelian period matches that of an Abelian power when $\boldsymbol{u}_0$ and $\boldsymbol{u}_k$ are both empty and $k > 2$.

As an example, the word $\boldsymbol{w} = abaababa$ over the alphabet $\Sigma = \{a, b\}$ can be written as $\boldsymbol{w} = \boldsymbol{u}_0\boldsymbol{u}_1\boldsymbol{u}_2\boldsymbol{u}_3$, where $\boldsymbol{u}_0 = ab$, $\boldsymbol{u}_1 = aab$, $\boldsymbol{u}_2 = aba$, $\boldsymbol{u}_3 = \varepsilon$, with ε the empty word, so that 3 is an Abelian period of $\boldsymbol{w}$ with Parikh vector $(2, 1)$ (the Parikh vector of $\boldsymbol{u}_0$ is $(1, 1)$ and that of $\boldsymbol{u}_3$ is $(0, 0)$ which are both "contained" in $(2, 1)$). Notice that $\boldsymbol{w}$ has also Abelian period 2, since it can be written as $\boldsymbol{w} = \boldsymbol{u}_0\boldsymbol{u}_1\boldsymbol{u}_2\boldsymbol{u}_3\boldsymbol{u}_4$, with $\boldsymbol{u}_0 = a$, $\boldsymbol{u}_1 = ba$, $\boldsymbol{u}_2 = ab$, $\boldsymbol{u}_3 = ab$, $\boldsymbol{u}_4 = a$.

This example shows that a word can have different Abelian periods. Moreover, a word can have the same Abelian period p corresponding to different factorizations; that is, with different heads. Actually, a word of length n can have $\Theta(n^2)$ many different Abelian periods [4], if these are represented in the form (h, p), where h is the length of the head—the length of the tail is uniquely determined by h and p.

Recently [5, 4] we described algorithms for computing all the Abelian periods of a word of length n in time $O(n^2 \times \sigma)$. This was improved to time $O(n^2)$ in [6]. In [7] the authors derived an efficient algorithm for computing the Abelian periods based on prior computation of the Abelian squares.

An Abelian period is called *full* if both the head and the tail are empty. Clearly, a full Abelian period is a divisor of the length of the word.

A preliminary version of the present paper appeared in [1] where we presented brute force algorithms to compute full Abelian periods and Abelian periods without head and with tail in $O(n^2)$ time and a quasi-linear time algorithm QLFAP for computing all the full Abelian periods of a word. In [8] Kociumaka et al. gave a linear time algorithm LFAP for the same problem. Here we first briefly outline LFAP, followed by a description of QLFAP. Then, extending the presentation in [1], we add an experimental section to demonstrate that our algorithm significantly outperforms LFAP in practice, both on pseudo-randomly generated and genomic data. Our method has the additional advantage of being conceptually simple and easy to implement.

2. Notation

Let $\Sigma = \{a_1, a_2, \ldots, a_\sigma\}$ be a finite ordered alphabet of cardinality σ and Σ^* the set of finite words over Σ. We let $|\boldsymbol{w}|$ denote the length of the word $\boldsymbol{w}$. Given a word $\boldsymbol{w} = \boldsymbol{w}[0\,..\,n-1]$ of length $n > 0$, we write $\boldsymbol{w}[i]$ for the $(i+1)$-th symbol of $\boldsymbol{w}$ and, for $0 \leqslant i \leqslant j < n$, we write $\boldsymbol{w}[i\,..\,j]$ for the factor of $\boldsymbol{w}$ from the $(i+1)$-th symbol to the $(j+1)$-th symbol, both included. We let $|\boldsymbol{w}|_a$ denote the number of occurrences of the symbol $a \in \Sigma$ in the word $\boldsymbol{w}$.

The *Parikh vector* of $\boldsymbol{w}$, denoted by $\mathcal{P}_{\boldsymbol{w}}$, counts the occurrences of each letter of Σ in $\boldsymbol{w}$, that is, $\mathcal{P}_{\boldsymbol{w}} = (|\boldsymbol{w}|_{a_1}, \ldots, |\boldsymbol{w}|_{a_\sigma})$. Notice that two words have the same Parikh vector if and only if one word is a permutation of the other (in other words, an anagram).

Given the Parikh vector $\mathcal{P}_{\boldsymbol{w}}$ of a word $\boldsymbol{w}$, we let $\mathcal{P}_{\boldsymbol{w}}[i]$ denote its i-th component and $|\mathcal{P}_{\boldsymbol{w}}|$ its norm, defined as the sum of its components. Thus, for $\boldsymbol{w} \in \Sigma^*$ and $1 \leqslant i \leqslant \sigma$, we have $\mathcal{P}_{\boldsymbol{w}}[i] = |\boldsymbol{w}|_{a_i}$ and $|\mathcal{P}_{\boldsymbol{w}}| = \sum_{i=1}^{\sigma} \mathcal{P}_{\boldsymbol{w}}[i] = |\boldsymbol{w}|$.

Finally, given two Parikh vectors $\mathcal{P}, \mathcal{Q}$, we write $\mathcal{P} \subset \mathcal{Q}$ if $\mathcal{P}[i] \leqslant \mathcal{Q}[i]$ for every $1 \leqslant i \leqslant \sigma$ and $|\mathcal{P}| < |\mathcal{Q}|$. This makes precise the notion of "contained" used in the Introduction.

Definition 1 (Abelian period [3]). A word $\boldsymbol{w}$ has an Abelian period (h, p) if $\boldsymbol{w} = \boldsymbol{u}_0\boldsymbol{u}_1 \cdots \boldsymbol{u}_{k-1}\boldsymbol{u}_k$ such that:

- $\mathcal{P}_{\boldsymbol{u}_0} \subset \mathcal{P}_{\boldsymbol{u}_1} = \cdots = \mathcal{P}_{\boldsymbol{u}_{k-1}} \supset \mathcal{P}_{\boldsymbol{u}_k}$,
- $|\boldsymbol{u}_0| = h$, $|\boldsymbol{u}_1| = p$.

We call $\boldsymbol{u}_0$ and $\boldsymbol{u}_k$ respectively the *head* and the *tail* of the Abelian period. Notice that the length $t = |\boldsymbol{u}_k|$ of the tail is uniquely determined by h, p and $|\boldsymbol{w}|$, namely $t = (|\boldsymbol{w}| - h) \bmod p$.

The following lemma gives a bound on the maximum number of Abelian periods of a word.

Lemma 1 ([5]). *The maximum number of different Abelian periods (h, p) for a word of length n over an alphabet of size σ is $\Theta(n^2)$.*

PROOF. The word $(a_1 a_2 \cdots a_\sigma)^{n/\sigma}$ has Abelian period (h, p) for any $p \equiv 0 \bmod \sigma$ and every h such that $0 \leqslant h \leqslant \min(p-1, n-p)$. □

An Abelian period is called *full* if it has head and tail both empty. We are interested in computing all the full Abelian periods of a word. Notice that a full Abelian period of a word of length n is a divisor of n. In the remainder of this note, we will therefore write that a word $\boldsymbol{w}$ has an Abelian period p if and only if it has full Abelian period $(0, p)$.

3. Previous work

We now outline the linear algorithm LFAP given in [8].

Let w be a word of length n. Let $\mathcal{P}_{w_i} = \mathcal{P}_{w[0..i]}$. Two positions $i, j \in \{1, \ldots, n\}$ are called proportional, which is denoted by $i \sim j$, if $\mathcal{P}_{w_i}[k] = c \times \mathcal{P}_{w_j}[k]$ for each k, where c is a real number independent of k.

An integer k is called a candidate (as a potential Abelian period) if $i \sim k$ for each $i \in \mathit{Mult}(k, n)$ where $\mathit{Mult}(k, n)$ is the set of multiples of k not exceeding n, or equivalently $k \sim 2k \sim 3k \sim \cdots$.

A positive integer $q \leq n$ is a full Abelian period of w if and only if $q \mid n$ and q is a candidate.

The algorithm first computes the set $A = \{k : k \sim n\}$ and then the set $B = \{d \in \mathit{Div}(n) : \mathit{Mult}(d, n) \subseteq A\}$.

Let ℓ be the rank in the alphabet of a least frequent letter a of w. Let q_0 be the position of the first occurrence of a in w. For $i \in \{q_0, \ldots, n\}$ let $\gamma_i = \mathcal{P}_{w_i} / \mathcal{P}_{w_i}[\ell]$. Vectors γ_i are used in order to deal with vector equality instead of vector proportionality.

If $i, j \in \{q_0, \ldots, n\}$ then $i \sim j$ is equivalent to $\gamma_i = \gamma_j$. The problem of computing the set A reduces to check if $\delta_i = \gamma_n - \gamma_i$ is equal to $(0, \ldots, 0)$ for $q_0 \leq i \leq n-1$. This is done in linear time by using what the authors called a diff-representation of the γ_i's (see [8] for details).

The authors show that the set B can be computed in linear time. This can be done by showing that after an $O(n)$ time preprocessing, the value $\gcd(k, \ell)$ for any $k, \ell \in \{1, \ldots, n\}$ can be computed in constant time (see [8] for details).

4. The new algorithm

In this section, we describe our algorithm QLFAP for computing all the full Abelian periods of a word $\boldsymbol{w}$ of length n over a constant-size alphabet $\Sigma = \{a_1, a_2, \ldots, a_\sigma\}$.

In a linear-time preprocessing phase, we compute $\mathcal{P}_{\boldsymbol{w}}[j]$ for $j = 1, 2, \ldots, \sigma$, i.e., the components of the Parikh vector of the word $\boldsymbol{w}$. Let g denote the greatest common divisor of the elements of the Parikh vector of $\boldsymbol{w}$, computable in $O(\sigma + \log n/\sigma)$ time [9]. Then we compute $s = n/g$. We can suppose $\sigma \geq 2$ and $g > 1$, otherwise the solution is trivial. In $O(\sqrt{g})$ time we compute a stack D of all divisors $d \geq 1$ of g in increasing order.

Observation 1. The only possible full Abelian periods of $\boldsymbol{w}$ are of the form $p = d \times s$, where d is an entry in D. Therefore, the smallest possible full Abelian period of $\boldsymbol{w}$ is s.

Definition 2 (scaled). A factor $\boldsymbol{u}$ of a word $\boldsymbol{w}$ is ***scaled*** if and only if there exists $k \geqslant 0$ such that $\mathcal{P}_{\boldsymbol{u}} = k \times \mathcal{P}_{\boldsymbol{w}}/s$, where $s = n/g$ is the smallest possible Abelian period of $\boldsymbol{w}$.

Observation 2. If $\boldsymbol{w} = \boldsymbol{uv}$ and $\boldsymbol{u}$ is scaled, then $\boldsymbol{v}$ is scaled.

A scaled factor is called *irreducible* if it cannot itself be factored into scaled factors. According to Definition 2 and Observation 2, every word $\boldsymbol{w}$ can be factored uniquely into irreducible scaled factors by computing the shortest scaled prefix $\boldsymbol{u}$ of $\boldsymbol{w} = \boldsymbol{uv}$ and then factoring the suffix $\boldsymbol{v}$ recursively, until this suffix is empty. This factorization is computed by calling procedure COMPUTEL given in Figure 1, which

```
ComputeL(w, s, g, P_w)
 1  (i, T, L) ← (0, 0, 0^|w|)
 2  while i ⩽ |w| − s do
 3     j ← 0
 4     t ← 0
 5     count ← 0^σ
 6     while j < σ do
 7        t ← t + 1
 8        for k ← 1 to s do
 9           count[w[i]] ← count[w[i]] + 1
10           i ← i + 1
11        j ← 0
12        while j < σ and ⌊count[j]/t⌋ = P_w[j]/g do
13           j ← j + 1
14     L[i − t × s] ← 1
15     T ← max(T, t)
16  return (L, T)
```

Figure 1: Algorithm computing array L such that $L[i] = 1$ iff i is the starting position of a scaled suffix of $\boldsymbol{w}$.

computes the scaled prefixes of $\boldsymbol{w}$ (or analogously, by the previous observation, the scaled suffixes of $\boldsymbol{w}$). It returns a boolean array L of length n defined by: $L[i] = 1$ if and only if i is the starting position of a scaled suffix of $\boldsymbol{w}$. It also returns the value T such that the longest scaled factor computed is of length $s \times T$.

Observation 3. Consider two positions i, j, with $0 \leqslant i < j < |\boldsymbol{w}|$, such that $L[i] = L[j] = 1$. Then $\boldsymbol{w}[i \,..\, j-1]$ is scaled.

Observation 4. A full Abelian period of w must be greater than or equal to $s \times T$.

Proposition 2. *The algorithm* ComputeL *computes the boolean array L of a word $\boldsymbol{w}$ of length n over an alphabet of size σ in time $O(n)$.*

Proof. The internal loop in line 6 is performed until a scaled factor of length $t \times s$ is found: the loop in line 8 computes the Parikh vector of the factor $w[i - t \times s \,..\, i]$ and the loop in line 12 verifies that this factor is scaled: its length is equal to $t \times s$ and it should hold for every letter j that $count[j]/(t \times s) = \mathcal{P}_{\boldsymbol{w}}[j]/n$. This can be rewritten as $count[j]/t = \mathcal{P}_{\boldsymbol{w}}[j]/g$ since $n = s \times g$. Since, by definition of g, $\mathcal{P}_{\boldsymbol{w}}[j]/g$ is an integer the test can be reduced to $\lfloor count[j]/t \rfloor = \mathcal{P}_{\boldsymbol{w}}[j]/g$ since it should hold for every letter. This avoids to use real numbers. Furthermore the $\mathcal{P}_{\boldsymbol{w}}[j]/g$ are constant for a given $\boldsymbol{w}$ and can thus be precomputed. When such a scaled factor is found, its starting position is $i - t \times s$ and $L[i - t \times s]$ is set to 1. Since the algorithm starts from position $i = 0$, therefore, according to Observation 2, array L is filled correctly.

The algorithm visits each position i in $\boldsymbol{w}$ exactly once, and corresponding to each i performs a constant-time processing. The internal loop in line 12 is performed at most σ times, every s positions. It is performed $|\boldsymbol{w}|/s$ times and since $\sigma \leqslant s$, each iteration costs $O(s)$ time. Thus the algorithm runs in time $O(|\boldsymbol{w}|)$. □

Example 1. $\boldsymbol{w}$ = abaababbbabaabbabbaaabbababbaa, $\mathcal{P}_{\boldsymbol{w}} = (15, 15)$.

i	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29
$\boldsymbol{w}[i]$	a	b	a	a	b	a	b	b	b	a	b	a	a	b	b	a	b	b	a	a	a	b	b	a	b	a	b	b	a	a
$L[i]$	1	0	1	0	0	0	0	0	1	0	1	0	1	0	1	0	1	0	0	0	1	0	1	0	1	0	1	0	0	0

The word $\boldsymbol{w}$ is factored into irreducible scaled factors in the following way: $\boldsymbol{w} = \mathtt{ab} \cdot \mathtt{aababb} \cdot \mathtt{ba} \cdot \mathtt{ba} \cdot \mathtt{ab} \cdot \mathtt{ba} \cdot \mathtt{bbaa} \cdot \mathtt{ab} \cdot \mathtt{ba} \cdot \mathtt{ba} \cdot \mathtt{bbaa}$. *Since the length of the longest word in this factorization is* 6 *and* $s = 2$, *the value of* T *is* 3.

Lemma 3. *The word $\boldsymbol{w}$ has full Abelian period p if and only if for every $k = 0, 1, \ldots, |\boldsymbol{w}|/p - 1$, $k \times p$ is the starting position of a scaled suffix of $\boldsymbol{w}$.*

Proof. Suppose that p is an Abelian period of $\boldsymbol{w}$. According to Observation 3, we have that for every $k = 0, 1, \ldots, |\boldsymbol{w}|/p - 1$ the position $k \times p$ is the starting position of a scaled factor and then of a scaled suffix of $\boldsymbol{w}$.

Conversely, if $k \times p$ is the starting position of a scaled suffix of $\boldsymbol{w}$ for every $k = 0, 1, \ldots, |\boldsymbol{w}|/p - 1$, it follows from Observation 3 that $\boldsymbol{w}[k \times p \,..\, (k+1) \times p - 1]$ is scaled and consequently that p is an Abelian period of $\boldsymbol{w}$. □

The algorithm QuasiLinearFullAbelianPeriods($\boldsymbol{w}$), given in Figure 2, computes all the full Abelian periods of a word $\boldsymbol{w}$. It first computes the greatest common divisor g of the elements of $\mathcal{P}_{\boldsymbol{w}}$ (line 1). Then it computes $s = n/g$ (line 2). This is followed by the computation of the array L giving the starting positions of the scaled suffixes and of the threshold T (line 3). Then it computes the stack D of all the divisors of g greater than or equal to T (line 4). Then it checks for every divisor d of g in D, whether all the multiples of d are starting positions of scaled suffixes (lines 6–12).

Example 2. $\boldsymbol{w} = \mathtt{abaababbbabaabbabbaaabbababbaa}$. *We have* $n = 30$, $\mathcal{P}_{\boldsymbol{w}}[1] = \mathcal{P}_{\boldsymbol{w}}[2] = 15$, $g = 15$, $s = 2$, $D = (1, 3, 5, 15)$. *When* $d = 1$, $p = 2$ *and* $4 = 2 \times p$ *is not a starting position of a scaled suffix. When* $d = 3$, $p = 6$ *and* 6 *is not a starting position of a scaled suffix. When* $d = 5$, $p = 10$ *and* 10, 20 *are starting positions of scaled suffixes, hence* 10 *is a full Abelian period of* $\boldsymbol{w}$. *The case where* $d = 15$ *is trivial since it corresponds to Abelian period* n. *Thus the algorithm returns* $\{10, 30\}$, *which is the set of full Abelian periods of* $\boldsymbol{w}$.

Theorem 4. *The algorithm* QuasiLinearFullAbelianPeriods($\boldsymbol{w}$) *computes all the full Abelian periods of a word $\boldsymbol{w}$ of length n over an alphabet of size σ in time $O(n \log\log n)$ and space $O(n)$.*

Proof. The correctness comes from Observation 1 and from Lemma 3. The algorithm QuasiLinearFullAbelianPeriods($\boldsymbol{w}$) scans all the multiples of the divisors $d \in D$, whose number is equal to the sum of the divisors of s, which in turn is $O(n \log\log n)$ [10]. The algorithm requires $O(n)$ space for L and D. □

5. Experimental results

5.1. Implementation

We implemented both QuasiLinearFullAbelianPeriods (QLFAP) and the algorithm LinearFullAbelianPeriods (LFAP) of [8] in a homogeneous way in C.

LFAP requires seven integer arrays of space $O(n)$ to compute the diff-representation plus five more for the preprocessing of the set B (see Sect. 3 for terminology). On the other hand, QLFAP requires only two integer arrays L and R of length n, together with the Parikh vector of length σ and the stack D (implemented as an integer array of length g).

In both algorithms, the memory space for the different arrays has been allocated one by one.

All the experiments have been performed on a computer with a 1.3 GHz Intel Core i5 processor and 4 Go 1600 MHz DDR3 RAM.

```
QuasiLinearFullAbelianPeriods(w = [w_0 .. w_{n-1}])
 1  g ← greatest common divisor of elements of P_w
 2  s ← n/g
 3  (L, T) ← ComputeL(w, s, g, P_w)
 4  D ← divisors of g greater or equal than T
 5  R ← ∅
 6  while d < g do
 7     k ← 1
 8     while k × d × s < n and L[k × d × s] = 1 do
 9        k ← k + 1
10     if n ⩽ k × d × s then
11        R ← R ∪ {i}
12     d ← Pop(D)
13  R ← R ∪ {n}
14  return R
```

Figure 2: Algorithm computing the set of Abelian periods of word $\boldsymbol{w}$.

5.2. Results

Figures 3 and 4 plots the raw execution times versus the word length (in logarithmic scale) while Figures 5–7 shows execution time ratios for LFAP over QLFAP. Experiments were performed on alphabet sizes 2, 5, 10 and 20, and input words of lengths 1000, 2000, ..., 10000. For each alphabet size and word length, 1000 words were randomly generated.

In order to verify the correct execution of both algorithms, words have been generated in such a way that they have at least one non trivial Abelian period. Both algorithms have been compiled with the same options and run on the same machine under the same conditions.

The plots show that, on the range of tested values, the two algorithms have a linear behaviour. The results show that our algorithm is significantly faster than LinearFullAbelianPeriods algorithm. One can also observe that the ratio increases both with word length and with the length of the Abelian period (see Figure 7). This may be due to the larger number of arrays needed by the LinearFullAbelianPeriods algorithm.

Similar results were obtained with the `E.coli` file from the Canterbury Large Corpus (Figures 8 and 9).

The source codes of LFAP and QLFAP algorithms can be found at `bioinfo.univ-rouen.fr/qlfap`.

6. Conclusions and perspectives

In this note we presented an algorithm for computing all the full Abelian periods of a word of length n in time $O(n \log\log n)$ and space $O(n)$. We showed that this algorithm is very efficient in practice, even if its theoretical worst-case behavior is superlinear.

Our experiments show that it is between 6 to 19 times faster than the linear time algorithm of [8] and we can expect that the ratio will be even larger for longer strings. Our method requires less space. In addition the latter algorithm needs to compare real numbers, which can be tricky in practice, while our method avoids this difficulty. Furthermore, our algorithm is conceptually simple, and stems from the notion of scaled factorization of a word, which could be of some interest in other related problems on abelian combinatorics on words.

A possible direction of further investigation could be that of extending the same approach to the computation of all the abelian periods with empty head but non-empty tail, or even to compute the smallest of such abelian periods.

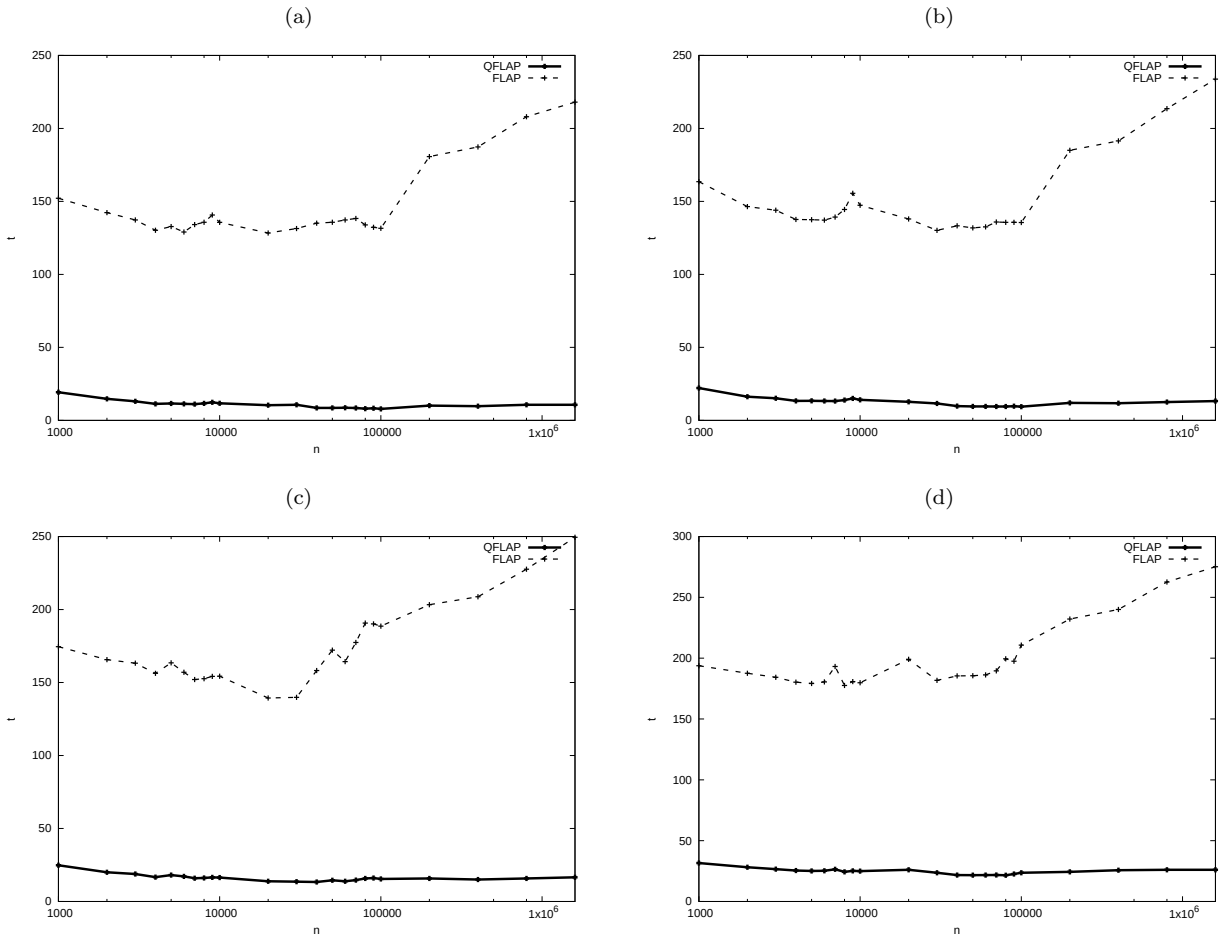

Figure 3: Running times of the LFAP and QLFAP algorithms for (a) alphabet of size 2 and at least one non trivial Abelian period 5; (b) alphabet of size 5 and at least one non trivial Abelian period 5; (c) alphabet of size 10 and at least one non trivial Abelian period 5; (d) alphabet of size 20 and at least one non trivial Abelian period 5.

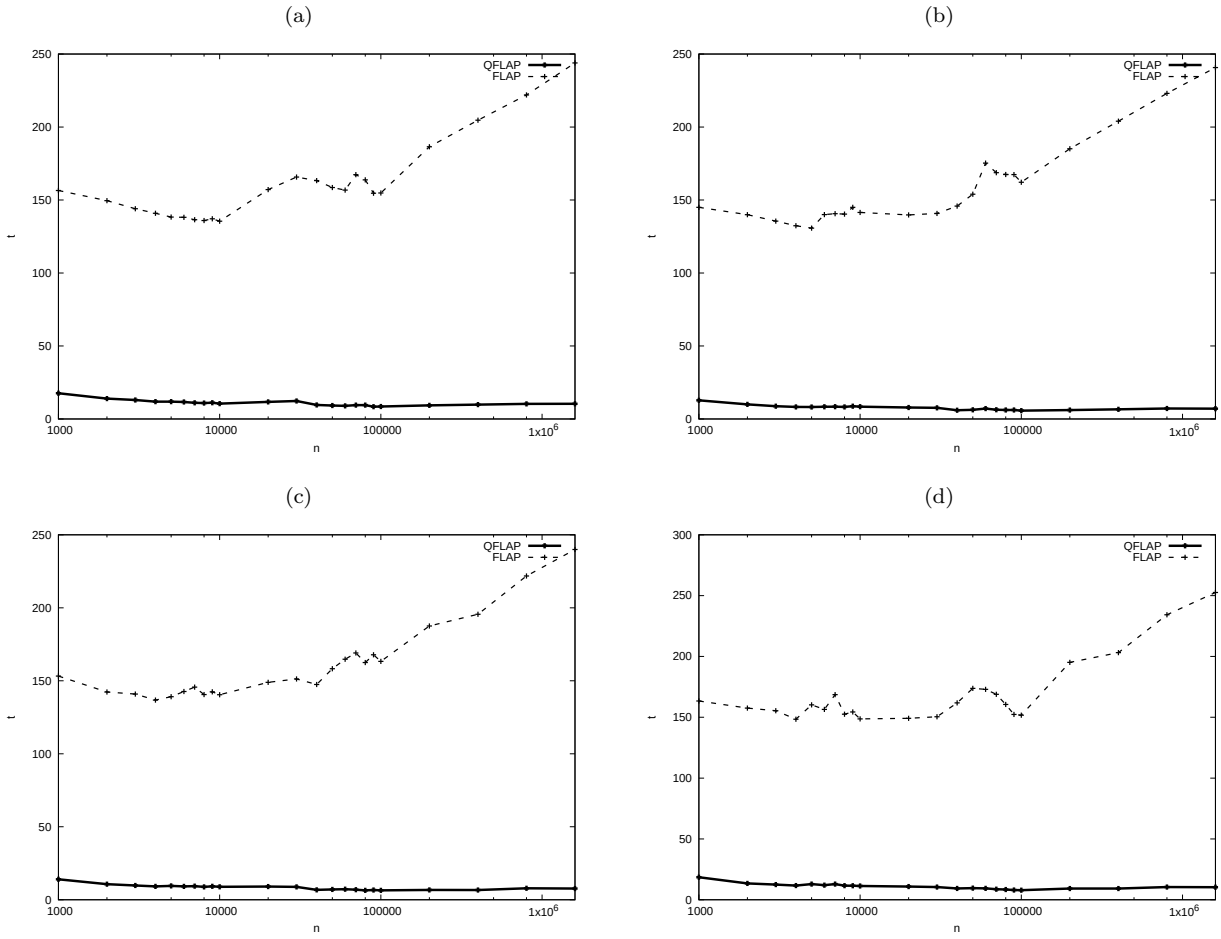

Figure 4: Running times of the LFAP and QLFAP algorithms for (a) alphabet of size 2 and at least one non trivial Abelian period 20; (b) alphabet of size 5 and at least one non trivial Abelian period 20; (c) alphabet of size 10 and at least one non trivial Abelian period 20; (d) alphabet of size 20 and at least one non trivial Abelian period 20.

(A) period 5

Alphabet	Word length 1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
2	7.90	9.64	10.51	11.50	11.45	11.42	12.02	11.64	11.38	11.62
5	7.38	9.03	9.50	10.32	10.26	10.32	10.54	10.38	10.34	10.46
10	7.04	8.28	8.67	9.37	9.03	9.13	9.55	9.47	9.33	9.47
20	6.13	6.66	6.93	7.08	7.14	7.16	7.31	7.30	7.18	7.20

(B) period 20

Alphabet	Word length 1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
2	8.88	10.72	11.06	11.88	11.67	11.87	12.32	12.48	12.32	12.90
5	11.33	13.98	15.47	15.99	15.91	16.70	16.71	16.98	16.57	16.77
10	10.93	13.30	14.45	14.95	14.69	15.57	15.62	15.90	15.44	16.03
20	8.81	11.72	12.46	12.74	12.46	13.06	13.07	13.17	13.35	13.59

Figure 5: Execution time ratios for LFAP over QLFAP algorithms. (A) Ratios computed on random words with at least one non trivial Abelian period 5. (B) Ratios computed on random words with at least one non trivial Abelian period 20.

(A) period 5

Alphabet	Word length 10000	20000	30000	40000	50000	60000	70000	80000	90000	100000
2	11.62	12.38	12.32	15.82	15.86	15.84	16.27	16.65	15.94	16.72
5	10.46	10.85	11.25	13.67	13.84	13.90	14.34	14.31	13.98	14.46
10	9.47	10.07	10.30	11.86	11.86	11.86	12.13	12.08	11.86	12.20
20	7.20	7.65	7.71	8.54	8.57	8.59	8.71	9.24	8.75	8.92

(B) period 20

Alphabet	Word length 10000	20000	30000	40000	50000	60000	70000	80000	90000	100000
2	12.90	13.43	13.44	17.00	17.23	17.46	17.63	17.21	18.37	18.10
5	16.77	17.80	18.45	24.27	24.58	24.69	26.55	27.07	27.20	28.42
10	16.03	16.50	17.17	21.80	22.46	22.85	24.47	25.31	25.09	25.45
20	13.59	13.73	14.37	17.41	18.27	18.49	19.43	19.07	19.02	19.33

Figure 6: Execution time ratios for LFAP over QLFAP algorithms. (A) Ratios computed on random words with at least one non trivial Abelian period 5. (B) Ratios computed on random words with at least one non trivial Abelian period 20.

(A) period 5

Alphabet	Word length 50000	100000	200000	400000	800000	1600000
2	15.86	16.72	17.87	19.29	19.50	20.35
5	13.84	14.42	15.44	16.24	16.99	17.66
10	11.86	12.20	12.88	13.89	14.44	15.08
20	8.54	8.92	9.53	9.34	10.07	10.57

(B) period 20

Alphabet	Word length 50000	100000	200000	400000	800000	1600000
2	17.23	18.10	20.04	20.82	21.49	23.39
5	24.58	29.42	30.21	30.76	31.16	34.13
10	22.46	25.45	27.84	29.42	28.33	31.23
20	18.27	19.33	21.17	22.05	22.42	24.45

Figure 7: Execution time ratios for LFAP over QLFAP algorithms. (A) Ratios computed on random words with at least one non trivial Abelian period 5. (B) Ratios computed on random words with at least one non trivial Abelian period 20.

(a)

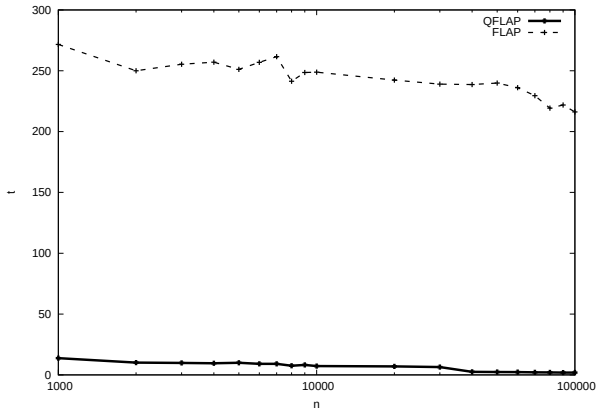

Figure 8: Running times of the LFAP and QLFAP algorithms for the `E.coli` file from the Canterbury Large Corpus.

Word length	1000	2000	3000	4000	5000	6000	7000	8000	9000	10000
Ratio	19.72	24.79	26.04	26.95	25.23	28.31	28.79	31.93	30.24	31.51

Word length	20000	30000	40000	50000	60000	70000	80000	90000	100000
Ratio	34.75	37.10	95.35	101.73	102.57	107.82	108.00	116.06	114.54

Figure 9: Execution time ratio between LFAP and QLFAP algorithms on words taken from the `E.coli` file from the Canterbury Large Corpus.

Acknowledgements

The authors thank the anonymous reviewers that greatly improved a first version of this paper.

References

[1] G. Fici, T. Lecroq, A. Lefebvre, E. Prieur-Gaston, W. F. Smyth, Quasi-linear time computation of the Abelian periods of a word, in: J. Holub, J. Zdárek (Eds.), Prague Stringology Conference, PSC 2012. Prague, Czech Republic, August 27–28, 2012. Proceedings, Czech Technical University in Prague, 2012, pp. 103–110.

[2] L. J. Cummings, W. F. Smyth, Weak repetitions in strings, Journal of Combinatorial Mathematics and Combinatorial Computing 24 (1997) 33–48.

[3] S. Constantinescu, L. Ilie, Fine and Wilf's theorem for abelian periods, Bulletin of the European Association for Theoretical Computer Science 89 (2006) 167–170.

[4] G. Fici, T. Lecroq, A. Lefebvre, É. Prieur-Gaston, Algorithms for computing Abelian periods of words, Discrete Applied Mathematics 163 (2014) 287–297.

[5] G. Fici, T. Lecroq, A. Lefebvre, É. Prieur-Gaston, Computing Abelian periods in words, in: J. Holub, J. Žďárek (Eds.), Proceedings of the Prague Stringology Conference 2011, Czech Technical University in Prague, Czech Republic, 2011, pp. 184–196.

[6] M. Christou, M. Crochemore, C. S. Iliopoulos, Identifying all abelian periods of a string in quadratic time and relevant problems, International Journal of Foundations of Computer Science 23 (6) (2012) 1371–1384.

[7] M. Crochemore, C. Iliopoulos, T. Kociumaka, M. Kubica, J. Pachocki, J. Radoszewski, W. Rytter, W. Tyczyński, T. Waleń, A note on efficient computation of all abelian periods in a string, Information Processing Letters 113 (3) (2013) 74–77.

[8] T. Kociumaka, J. Radoszewski, W. Rytter, Fast algorithms for abelian periods in words and greatest common divisor queries, in: Proc. of STACS, 2013, pp. 245–256, available from http://stacs-conf.org.

[9] G. Bradley, Algorithm and bound for the greatest common divisor of n integers, Comm. ACM 13 (7) (1970) 433–436.

[10] T. H. Gronwall, Some asymptotic expressions in the theory of numbers, Transactions of the American Mathematical Society 14 (1913) 113–122.